determining a synchronization set by:

(i) determining which, if any, information records have been previously transmitted to the target dataset but no longer exists at the source dataset, and

(ii) determining which, if any, information records have been added to or modified at the source dataset since the source dataset was last synchronized with the target dataset,

wherein each information record of the source dataset is assigned a globally unique identifier that is independent of either of the devices, for identifying said each information record at both the source dataset and the target dataset, said globally unique identifier being maintained in a device-independent record map that allows the globally unique identifier to be traced back to a specific information record regardless of which device the specific information record resides; and

based on said synchronization set, synchronizing information records of the source dataset with information records of the target dataset by:

(i) using said globally unique identifiers, deleting from the target dataset any information records which have been previously transmitted to the target dataset but no longer exist at the source dataset, and

(ii) using said globally unique identifiers, updating the target dataset so that said target dataset includes those information records determined to have been added to or modified at the source dataset since the source dataset was last synchronized with the target dataset.

---

2. The method of claim 1, wherein each dataset comprises a database table having a plurality of data records.

3. The method of claim 1, wherein each dataset comprises an electronic address book listing contact information.

4. The method of claim 1, wherein each dataset comprises an electronic schedule listing scheduling information.

5. The method of claim 1, wherein said globally unique identifiers are created by the system regardless of whether the source dataset includes existing record identifiers.

6. The method of claim 5, wherein said globally unique identifiers are maintained in a record map stored apart from the source dataset.

7. The method of claim 1, wherein each said globally unique identifier for each record comprises a timestamp which is assigned to the record when the record is initially processed by the system.

8. The method of claim 1, wherein each globally unique identifier is a 32-bit value.

9. The method of claim 1, further comprising:

synchronizing the information records of the target dataset with information records of the source dataset by designating the source dataset as the target dataset, designating the target dataset as the source dataset, and repeating said determining step and said synchronizing step.

10. The method of claim 1, wherein said synchronization set comprises a delete order specifying particular information records to delete at the target dataset.

11. The method of claim 10, wherein said delete order includes a list of globally unique identifiers for particular information records to delete at the target dataset.

12. The method of claim 1, wherein said synchronization set comprises an extraction record specifying particular information to add to or modify at the target dataset.

13. The method of claim 12, wherein said extraction record includes at least one globally unique identifier together with field information for the particular information to add to or modify at the target dataset.

14. The method of claim 1, further comprising:

excluding certain information records from participating in synchronization by applying a user-defined filter.

15. The method of claim 14, wherein said user-defined filter comprises an outbound filter applied to information records prior to creation of the synchronization set.

16. The method of claim 14, wherein said user-defined filter comprises an inbound filter applied to information records after creation of the synchronization set.

17. The method of claim 14, wherein said user-defined filter comprises a user-supplied filtering routine supplying filtering logic.

18. The method of claim 1, wherein said target dataset resides at a remote location relative to the source dataset.

19. The method of claim 18, further comprising:

after creating the synchronization set, transmitting said synchronization set to said remote location.

20. The method of claim 19, wherein the synchronization set is transmitted to the remote location using an electronic messaging communication protocol.

21. (Amended) A synchronization system comprising:

means for connecting a first device having a first dataset to a second device having a second dataset;

means for determining information of said first and second datasets which requires synchronization, said means including:

(i) means for determining for each dataset information which has been previously received from the other dataset but which no longer exists at the other dataset, and

(ii) means for determining for each dataset information which has been added or modified at the other dataset since the other dataset was last synchronized with said each dataset; and

means, responsive to said determining means, for synchronizing said first and second datasets;

wherein said information of said first and second datasets comprises data records and wherein said means for determining include means for assigning to each data record a device-independent globally unique identifier created by the system for uniquely identifying each data record regardless of which dataset and device it appears.

22. The system of claim 21, wherein at least one of said devices is a hand-held computing device.

23. The system of claim 21, wherein at least one of said devices is desktop computing device.

24. The system of claim 21, wherein said means for connecting includes a Transmission Control Protocol/ Internet Protocol (TCP/IP) connection.

25. The system of claim 21, wherein said means for synchronizing operates to provide bi-directional synchronization of the datasets.

26. --CANCELED--

27. The system of claim 21, further comprising:
filter means for selectively blocking synchronization of certain types of information.

28. The system of claim 27, wherein said filter means operates based on how old information is.

29. The system of claim 27, wherein said filter means operates based on particular information content.

30. The system of claim 21, further comprising:
electronic mail transport means for enabling synchronization of remote datasets.

Remarks

General

A. Status of application

Claims 1-30 have been examined and stand rejected in view of prior art. The claims have been amended in an effort to better distinguish the claimed invention. Reexamination and reconsideration of the claims, as amended, are respectfully requested.

B. Confirmation of commonly-owned subject matter

Applicants confirm that the subject matter of the various claims was commonly owned at the time of the invention(s) covered therein.

The invention

Applicants have invented a synchronization system with methods for synchronizing information among disparate datasets. In basic operation, the core engine of the synchronization system issues requests to accessors components which, in turn, instruct the relevant device (i.e., data-storing/processing device) to perform an appropriate data item or record operation, such as inserting a new record, or updating or deleting an existing record. The core engine asks the accessors (operating in conjunction with a conduit component) to provide those records that have changed since the last synchronization. Each such record is provided in the form of an extract record. Each extract record is processed according to the outbound synchronization logic and, if it has not been removed, is then mapped to a record map which provides a global identifier or ID and a timestamp.

Global IDs, which are employed at the level of a Record Map, are global to the entire synchronization system. Even if a source dataset already provides unique IDs, those unique IDs are generally unique to the device or unique to particular records on that device. Here, the synchronization system provides a unique global identifier (e.g., 32-bit or 64-bit value) for each data item at the level of the Record Map. Each global identifier can be based on an existing identifier, such as an existing record ID (RID) value, or can be synthesized de novo at runtime (e.g., based on system time/date information), particularly in those instances where the underlying dataset does not provide any identifiers for its data items. Regardless of how a global ID is provided for a given data item, that global ID is employed throughout the synchronization process for supporting synchronization across multiple devices (e.g., from palmtop to desktop to server) without creating duplicates. In conjunction with the system storing global IDs for identifying particular data items, the system also tracks when each data item or record is last modified (or created) by storing a last-modified date stamp. The actual change which occurred is logged in the Record Map.

Internally, therefore, the identification of each record is tracked together with information about when each record was last changed, so that proper synchronization can be carried out across multiple devices.

The basic flow or propagation of a record change from one dataset to another is as follows. At the outset, the changed record is extracted and looked up in the record map. The addition of the new record is noted in a Transaction Table, so the system can determine whether the record has already been dealt with. A corresponding export record (i.e., record data along with a globally unique ID) is generated for inclusion in the synchronization set or SyncSet being created. The SyncSet contains all of the record actions that need to be transmitted to the other device, including insertions, updates, or deletions. The SyncSet can, at this point, be sent via a variety of transport mechanisms, including e-mail, FTP (file transport protocol), and two-way pagers. The receiving device processes the SyncSet by proceeding to map it, during inbound synchronization. Here, the corresponding import record is mapped into a record form for the target dataset. Once the appropriate records get inserted, updated, or deleted, the inbound Transaction Table is updated. In this fashion, the present invention provides a synchronization system and methodology that allows each device to remain synchronized with all other devices in a convenient, transparent manner.

Prior art rejections

A. Rejection under 35 U.S.C. Section 102(e)

Claims 21-23, 25, and 27-29 stand rejected under 35 U.S.C. Section 102(e) as being anticipated by Meyering, U.S. Patent Number 5,729,735 (hereafter, "Meyering"). The Examiner equates the claims as reading on the relatively simplistic synchronization approach of Meyering, which describes a synchronization technique using backup files to reflect contents of a remote file at a particular point in time (i.e., when it was created or last synchronized). Actual synchronization involves a simple comparison between the master file, the remote file, and the backup file, to determine which file (i.e., remote or master) has a more current version of the data. Applicants' claims that were rejected under Section 102(e) have been amended to prevent such an interpretation.

In particular, independent claim 21 has been amended to include the global identifier limitation set forth in claim 26. Therefore, as amended, claim 21 includes the limitation:

> wherein said information of said first and second datasets comprises data records and wherein said means for determining include means for *assigning to each data record a device-independent globally unique identifier created by the system for uniquely identifying each data record regardless of which dataset and device it appears.* (emphasis added)

It is believed that the limitation clearly distinguishes the claims from that described by Meyering. Moreover, in the Section 102(e) rejection, the Examiner purposefully did not include claim 26 in the rejection, thus indicating that the Examiner also believed that the limitation distinguishes over the reference. Accordingly, it is respectfully submitted that the amendment distinguishes the claims over Meyering and that the rejection under Section 102(e) is overcome.

B. First rejection under 35 U.S.C. Section 103

Claims 1-13 stand rejected under 35 U.S.C. Section 103(a) as being unpatentable over Kucala, U.S. Patent Number 5,727,202 (hereinafter, "Kucala") in view of Olds et al., U.S. Patent Number 5,832,487 (hereinafter, "Olds"). The Examiner's rejection of claim 1 is representative:

> As to Claim 1, Kucala teaches a system which including 'a system providing one dataset in communication with another dataset, a method for synchronizing datasets' [col 1, line 8-10], 'receiving a request specifying synchronization of information records of a source dataset with information records of a target dataset' [fig 1, col 2, line 66-67, col 3, line 1-9], examiner interpreting source dataset to be equivalent to Kucala's fig 1, element 101, target dataset to bc equivalent to Kucala's fig 1, element 201, 'determining which, if any, information records have been previously transmitted to the target dataset but no longer exists at the source dataset' [col 3' line 6-9, line 35-53], 'determining which, if any, information records have been added to or modified at the source dataset since the source dataset was last synchronized with the target dataset' [col 3, line 6-9, line 35-63], 'synchronizing information

> records of the source dataset with information records of the target dataset' [col 3, line 54-63].

However, the Examiner acknowledges that Kucala does not teach the "globally unique identifier" limitation of the claims. To remedy this deficiency, the Examiner turns to Olds as teaching this limitation:

> Olds teaches a system which including 'source dataset' [col 3, line 20-24], examiner interpreting source dataset to be equivalent to Olds's file servers, element 16, ' assigned a globally unique identifier for identifying said each information record, [col 3, line 25-30], 'at both source dataset and the target dataset' [col 3' line 42-45], examiner interpreting target dataset to be equivalent to Olds's Clients, element 20, 'globally unique identifiers' [col 3, line 54-57], deleting from the target dataset any information records which have been previously transmitted to the target dataset but no longer exist at the source dataset'[col 10, line 7-14, line 19-28], 'using said globally unique identifiers, updating the target dataset so that said target dataset includes those information records determined to have added to or modified at the source dataset since the source dataset was last synchronized with the target dataset' [col 10, line 15-28].

Therefore, the Examiner concludes that it would have been obvious to one of ordinary skill in the art at the time of the Applicant's invention to combine Olds with Kucala "because globally unique identifiers allows [one] to identify particular dataset or record to be update or modify, such an identifier is unique because it also contains a time stamp, eliminates duplicate records, chances of confusion over a conflict are greatly reduced, save time in later accesses or reorganization of that datasets or records, thus improving the reliability and versatility of the system." A detailed review of the references, however, reveals that Applicants' claimed invention may be distinguished on a variety of grounds.

Kucala describes a fairly simplistic synchronization approach that uses backup files to reflect the status (e.g., new, updated, or deleted) of records on a palmtop device (e.g., Palm handheld computing device), as well as to reflect the status of records on a desktop device (e.g., PC computer). The results of both compares, which are stored in a temporary data structure (e.g., a "reconcile file"), are copied over the selected files on the palmtop, the

PC, and the backup file. Thus the approach described by Kucala is similar to that described by Meyering (above).

Olds describes an alternative approach to achieving the functionality provided by GUIDs (globally-unique IDs), such as allowing identification of an object independently of names imposed by users. The approach operates within a system that includes a hierarchical set of objects where any contiguous subtree of objects may be declared to be a "partition." A partition may exist at more than one location. A partition at a particular location is called a "replica" and is assigned a unique number among the set of replicas of its partition. Each server holds at most one replica of a given partition. When an object is created within a replica it is assigned a unique identifier relative to its sibling objects, such that no two siblings in any replica of the partition have the same relative identifier. The relative identifier contains its replica number and is therefore guaranteed to be unique among all of the sibling objects created in all replicas of that partition. The approach provides an object identifier for an object by starting with the relative identifier of an object and adding the relative identifier of each ancestor object until the relative identifier of the root of the object hierarchy is added. Such a sequence of relative object identifiers can be used as a global identifier of an object and is referred as a "tuned-name" or an "absolute name." Olds distinguishes his identifiers from globally-unique IDs. For instance, Olds states that his tuned-names, unlike GUIDs, support arbitrary levels of object hierarchy both within the total system and within a replica. The complete sequence of relative identifiers is significant. Moreover, GUIDs may collide. That is, two objects might be identified by the same GUID. By contrast, tuned-names are unique. Additionally, Olds states that an object identified by a tuned-name can be located in a step-by-step fashion starting from any point the object hierarchy, while GUIDs, by contrast, have no internal structure that can be used to narrow the search for an object from all possible locations. Clearly, Olds seeks to distinguish his approach from globally-unique identifiers. Although his approach may prove to be a useful alternative to globally-unique identifiers, his approach in no way remedies the deficiencies of Kucala.

Applicants' claimed invention does not seek to replace globally-unique identifiers or GUIDs with an alternative identifier, such as Olds' "tuned-names." The GUIDs per se that are employed in Applicants' system may in fact be rather conventional identifiers, such as a 32-bit or 64-bit value quantity (e.g., integer). It is not so much the identifiers themselves that are of interest but how they are used during the synchronization process. During operation of the present invention, the participating data sets (i.e., source and target data sets) typically will already have their own GUID values (e.g., assigned by their own computing device). In the claimed invention, the GUIDs -- whether the aforementioned device-provided GUID values or newly synthesized values -- are employed internally by the synchronization system (e.g., at the level of the synchronization system's Record Map) in a manner that allows each data item to be uniquely identify across multiple disparate devices, regardless of whether such devices may or may not be connected to a network and regardless of whether such devices have identification schemes that are compatible with one another. By tracking the identification of each record in a manner which is independent of what device or system the record resides and doing so at the level of the synchronization's system record map, Applicants' invention allows proper synchronization to be carried out across multiple devices in a convenient, transparent manner.

Although the originally-filed claims are believed to be distinguishable over the cited references, the claims have nevertheless been amended in an effort to further distinguish Applicants' invention. For instance, claim 1 has been amended to include the limitation (shown in amended form):

> wherein each information record of the source dataset is assigned a globally unique identifier that is independent of either of the devices, for identifying said each information record at both the source dataset and the target dataset, said globally unique identifier being maintained in a device-independent record map that allows the globally unique identifier to be traced back to a specific information record regardless of which device the specific information record resides; [...]

The amendment to the claim further clarifies how the globally unique identifier is employed. It is respectfully submitted that amended claim 1 (as well as its dependents) distinguishes

Applicants' invention from the cited references and that the rejection of claims 1-13 under Section 103 is overcome.

C. Second rejection under 35 U.S.C. Section 103

Claims 14-17 stand rejected under 35 U.S.C. Section 103(a) as being unpatentable over Kucala and Olds as applied to claim 1 above, and further in view of Buchanan, U.S. Patent Number 5,758,355 (hereinafter, "Buchanan"). Here, the Examiner cites Buchanan for the additional teaching of a filter limitation. The claims are believed to be allowable for at least the reasons cited above pertaining to Kucala and Olds.

Additionally, the claims are believed to be allowable for the following additional reasons. In addition to including the limitation of a "filter," the claims specify how a specific filter is applied in regards to the synchronization process. For instance, claim 15 specifies that the filter "comprises an outbound filter applied to information records prior to creation of the synchronization set." Claim 16, on the other hand, specifies that the filter "comprises an inbound filter applied to information records after creation of the synchronization set." Even if the Examiner interprets the "filter" limitation to be equivalent to selection criteria in a database query (e.g., SQL predicate), the Buchanan reference still fails to teach or suggest specific application of such a "filter" in relation to creation of a synchronization set, as required by Applicants' claims. Accordingly, it is submitted that the claims distinguished over the cited references and that the rejection, particularly in light of clarifying amendments made to parent claim 1, is overcome.

D. Third rejection under 35 U.S.C. Section 103

Claims 24, 26, and 30 stand rejected under 35 U.S.C. Section 103(a) as being unpatentable over Meyering as applied to claim 21 above, and further in view of Olds. Claim 26 has been canceled, pursuant to previously-discussed amendments to the claims. The remaining claims are believed to be allowable for at least the reasons cited above pertaining to Meyering and the reasons cited above pertaining to Olds.